

B<sup>c</sup> destination site 204 is able to identify which columns are held in common in the propagated changes by performing an intersection of the destination flavor 234 and the source flavor 232 to generate a cursor that implements that intersection when invoked to insert or update rows at the master tables (step 320).

---

Page 14, last paragraph, lines 21-25, substitute therefor:

---

B<sup>2</sup> Suppose the developers of the database application wish to evolve the schema by adding a new column, B, to replicated table 410 and dropping the column A 422, but with little disruption to the entire distributed database system. Accordingly, the object group for the replicated tables 410, 411, and 412 is given a flavor KA 430 that states that the replicated tables 410, 411, and 412 has a primary key 420 column and column A 422.

---

Page 15, first paragraph, lines 1-115, substitute therefor:

---

B<sup>3</sup> FIG. 5 shows steps taken in the developers' phased roll-out. At step 500 and with concurrent reference to FIG. 4(b), the developers define a new top flavor KAB 432 at the headquarters site 400 that includes the new column B 424 as well as the obsolescent column A 422. Flavor KAB 432 is called the "top" flavor because it describes the fullest extent of the object group. The definition for the flavor KAB 432 is published to the other sales office sites 402 and 404, before instantiating the flavor KAB 432 at the headquarters site 400 so that the sales office sites 402 and 404 will know how to handle propagated changes coming from the headquarters site 400 in the new flavor KAB 432. Instantiating the flavor KAB 432 at the headquarters site 400 causes the new column 424 to be added to old table 410 to produce new table 413 with a different shape than previously.

---